



Intel® E8500 Chipset North Bridge (NB)

Specification Update

June 2005

Notice: The Intel® E8500 chipset North Bridge (NB) may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are documented in this specification update.

Document Number: 306747-002



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® E8500 chipset North Bridge (NB) may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel, Pentium, Intel Xeon, Intel NetBurst, Intel SpeedStep, Intel Extended Memory 64 Technology and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2005, Intel Corporation



Contents

Preface 5

Summary Table of Changes..... 6

Identification Information 8

Errata..... 9

Specification Changes..... 15

Specification Clarifications 16

Documentation Changes..... 18

Revision History

Version	Description	Date
-001	Initial Public Release	March 2005
-002	Added errata 15-16; added Specification Clarification 3.	June 2005

Preface

This is an update to the specifications in the documents listed in the “[Affected Documents](#)” and “[Nomenclature](#)” tables. It is a compilation of device and document errata and specification clarifications/changes, and is intended for hardware system manufacturers and software developers.

Information types defined in the [Nomenclature](#) section of this document are consolidated into this document and are no longer published in other documents. This document may also contain previously unpublished information.

Affected Documents

Document Title	Document Number
Intel® E8500 Chipset North Bridge (NB) Datasheet	306745

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L3 cache size, package type, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

QDF Number is a several digit code used to distinguish between engineering samples. These samples are used for qualification and early design validation. The functionality of these parts can range from mechanical only to fully functional. This document has a processor identification information table that lists these QDF numbers and the corresponding product sample details.

Errata are design defects or errors. These may cause the Short Product Name behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Changes/Clarifications are modifications to the current published specifications. These changes will be incorporated in the next release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in the next release of the specification.

Note: Errata remain in the specification update throughout the product's life cycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update are archived and available upon request. Specification changes, specification clarifications, and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation.



Summary Table of Changes

The tables included in this section indicate the errata, specification changes/ clarifications, or documentation changes that apply to the Intel® E8500 chipset North Bridge (NB). Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted.

Codes Used in Summary Table


Stepping/Version

- X: Applies to this stepping.
- (No mark) or (Blank box): Fixed in listed stepping or does not exist in listed stepping.

Status

- Plan Fix Root cased to a silicon issue and will be fixed in a future stepping.
- Fixed Root cased to a silicon issue and has been fixed in a subsequent stepping.
- No Fix Root caused to a silicon issue that will not be fixed.

Row

-  Change bar to left of table row indicates that this item is either new or modified from the previous version of this document.

Errata

Number	Stepping	Status	ERRATA
	B-0		
1	X	No Fix	Hub interface data parity error may not generate poison
2	X	No Fix	FSB ECC error followed by a 0-length transfer logs two errors
3	X	No Fix	Uncorrectable errors are logged twice
4	X	No Fix	VPP ports may report errors across system resets
5	X	No Fix	PCI Express* Errors should not be routed to H_{A,B}_MCERR#
6	X	No Fix	PCI Express Cards Intermittently Train x4 Instead of x8
7	X	No Fix	PCI Express Port A0 will not Train x1
8	X	No Fix	NB may hang under heavy stress
9	X	No Fix	EXP_SLOTSTS.CMDCOMP cannot be cleared
10	X	No Fix	Replay Timer errors logged when PCI Express ports exit L1 power management state
11	X	No Fix	NB hang during stress and PCI Express link events
12	X	No Fix	NB may cause IERR# when using SYRE.CPURST with the Intel® Xeon™ Processor MP with 1 MB L2 Cache
13	X	No Fix	NB may trigger ECC errors during warm reset
14	X	No Fix	NB PCI Express ports get stuck in reset state
15	X	No Fix	Receiver errors reported on NB when link is disabled
16	X	No Fix	Receiver errors reported on Port 4 when link is disabled in Port 6 in the 2 x8 configuration

Specification Changes

Number	SPECIFICATION CHANGES
NA	There are no Specification Changes in this revision of the Specification Update.

Specification Clarifications

Number	SPECIFICATION CHANGES
1	Error detection for PCI Express devices not present should be disabled
2	PCI Express links enter compliance mode if connected to a powered down device
3	PCI Express: Configuration completion time-out enable needs non-configuration completion time-out to be enabled in order to work

Documentation Changes

Number	DOCUMENTATION CHANGES
NA	There are no Documentation Changes in this revision of the Specification Update.

Identification Information

Component Identification via Programming Interface

The Intel® E8500 chipset North Bridge (NB) can be identified by the following register contents:

Stepping	Vendor ID ^a	Device ID ^b	Revision Number ^c
B-0	8086h	2600h	01h

NOTES:

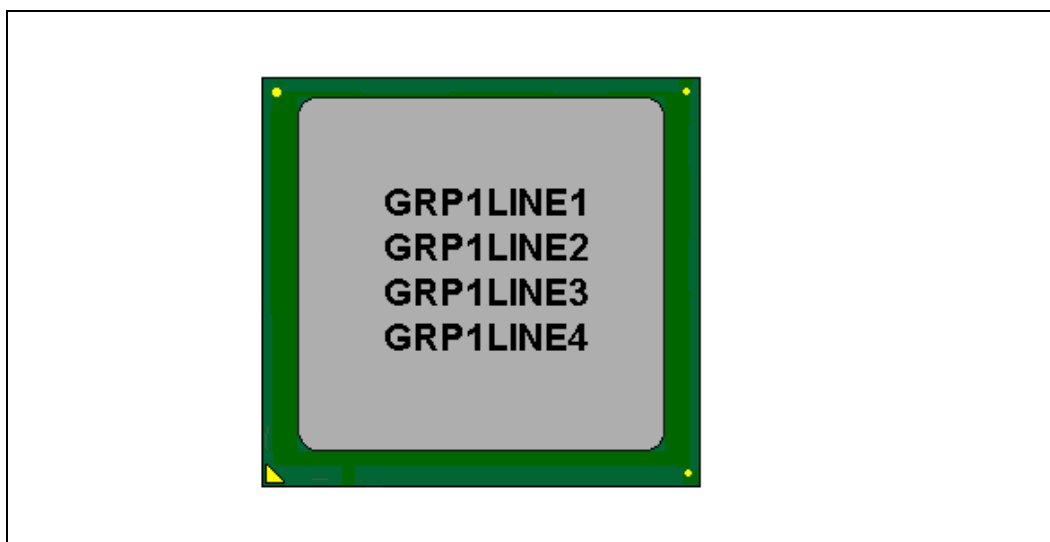
- The Vendor ID corresponds to bits 15:0 of the Vendor ID Register located at offset 00 - 01h in the PCI function 0 configuration space.
- The Device ID corresponds to bits 15:0 of the Device ID Register located at offset 02 - 03h in the PCI function 0 configuration space.
- The Revision Number corresponds to bits 7:0 of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

Component Marking Information

The Intel® E8500 chipset North Bridge (NB) stepping can be identified by the following component markings:

Stepping	QDF-Spec	S-Spec	Top Marking	Notes
B0	QE34ES	N/A	NQ84001TNB	
B0	QG54	N/A	QG84001TNB	Lead free package

Figure 1. Top-Side Marking Example (optional)



Errata

1. Hub interface data parity error may not generate poison

Problem: If a data parity error occurs in the very last data double word of an incoming hub interface transaction and the transaction size is a multiple of 4 double words, the poison information does not get propagated to CDC.

Implication: If this error is masked, silent data corruption may occur.

Workaround: Hub interface parity errors must be treated as fatal to contain this possible data corruption.

Status: For the steppings affected, see the *Summary Table of Changes*.

2. FSB ECC error followed by a 0-length transfer logs two errors

Problem: The NB does an ECC check on the data from the previous transaction on a 0-length transfer.

Implication: If a data ECC error occurs on the FSB subsequently followed by a 0 length transfer, the NB will log the ECC error twice.

Workaround: None.

Status: For the steppings affected, see the *Summary Table of Changes*.

3. Uncorrectable errors are logged twice

Problem: If an uncorrectable error occurs, this error will be logged in both the NB's FERR and NERR registers. This issue does not affect errors on IMI.

Implication: Details on errors occurring before an uncorrectable error will be lost.

Workaround: None.

Status: For the steppings affected, see the *Summary Table of Changes*.

4. VPP ports may report errors across system resets

Problem: VPP Ports may report errors in the EXP_unitERR.T9Err bit across system resets. Per the I²C* specification a data line can change when the clock is low. The only exceptions to this are Start and Stop commands in which is a master drives a data falling edge when clock is high (for Start) or drives a data rising edge when clock is high (for Stop).

The NB VPP state machine is constantly cycling through each VPP port and reading its state. If the NB gets reset when the clock is high and one of the VPP ports is driving the data line low (as part of the read or an acknowledge) then, upon reset, the VPP will not recognize the initial Start command and will not properly initialize. The NB recognizes the VPP did not respond properly and sets the T9Err bit.

Implication: Hot plug VPPs may non-deterministically be unresponsive after a reboot.

Workaround: This issue can be avoided by reinitializing the VPP ports when this error occurs as follows:

- Disable Error Logging (EXP_unitDMASK.T9DetMsk) for this port
- Enable this VPP
- Delay 50 ms

- Disable this VPP
- Wait 1 second
- Clear the Error bit (EXP_unitERR.T9Err)
- Enable Error Logging (EXP_unitDMASK.T9DetMsk)
- Enable VPP

Status: For the steppings affected, see the *Summary Table of Changes*.

5. PCI Express* Errors should not be routed to H_{A,B}_MCERR#

Problem: If PCI Express* errors are routed to the FSB H_{A,B}_MCERR# pins via settings in the EXP_ERR_DOCMD register, when an error occurs, MCERR# will be continually asserted (3 cycles on, 3 cycles off) until the error is cleared.

Implication: MCERR# should only be asserted for two cycles per error. Continuous assertion of MCERR# will result in an interrupt storm and system hang.

Workaround: PCI Express errors routed via the EXP_ERR_DOCMD register (D1-7, F0, 0x148) should be routed through one of the ERR[2:0] signals rather than the MCERR# pins.

Status: For the steppings affected, see the *Summary Table of Changes*.

6. PCI Express Cards Intermittently Train x4 Instead of x8

Problem: Intel validation boards which support x8 links will intermittently train as x4 links coming out of an AC power cycle only. This behavior seems to follow cards rather than systems. Some cards show this behavior up to 25% of the time.

Implication: This behavior has not been reproduced on other cards. Cards will not experience full bandwidth if they train as a x4 rather than x8 link.

Workaround: All ports should be strapped to their maximum widths (i.e. x8 slots should be strapped as x8, regardless of the populated card width) via the EXP_WIDTH strapping options.

Status: For the steppings affected, see the *Summary Table of Changes*.

7. PCI Express Port A0 will not Train x1

Problem: A logic error in NB B0 occasionally prevents port A0 lane 0 from gaining symbol lock.

Implication: During a x1 negotiation (i.e. with a x1 device or after x4 link training has failed because of a bad lane), link training will fail.

Workaround: A workaround for this bug is described in the BIOS Spec Update. This workaround must be run each time the link trains with a x1 device.

Status: For the steppings affected, see the *Summary Table of Changes*.

8. NB may hang under heavy stress

Problem: Under heavy memory write and PCI configuration space stress, internal NB's queues may become completely full and enter an arbitration deadlock

Implication: Under these conditions, the NB may hang the system.

Workaround: A workaround for this issue which adjusts internal buffer arbitration policy is described in the BIOS Spec Update.

Status: For the steppings affected, see the *Summary Table of Changes*.

9. **EXP_SLOTSTS.CMDCOMP cannot be cleared**

Problem: Due to the width of the internal NB bus, a write to clear the EXP_SLOTSTS.CMDCOMP bit causes the NB to also register a write to the EXP_SLOTCTRL register. As a result, the NB will set the EXP_SLOTSTS.CMDCOMP within a few cycles.

Implication: The EXP_SLOTSTS.CMDCOMP bit is effectively stuck high

Workaround: BIOS and OS's should not rely on the value of this bit and disable command complete interrupts by clearing setting EXP_SLOTCTRL.CCIEN

Status: For the steppings affected, see the *Summary Table of Changes*.

10. **Replay Timer errors logged when PCI Express ports exit L1 power management state**

Problem: When a PCI Express link is in L1, the link is in electrical idle, but the NB physical layer incorrectly indicates that it is ready to accept transactions to the NB link layer. If the NB issues a transaction during this time, it is dropped by downstream device, but does cause an exit from L1. The link goes into recovery and then L0. Since the initial transaction was dropped, there is no ack sent by the downstream device, which eventually causes the replay timer to expire.

Implication: A replay timer error may be logged whenever a PCI Express link exits L1. This is a correctable error. The dropped transaction will be resent and the link will operate normally after this error.

Workaround: COREDMASK.IO11DETMASK should be set to disable logging of replay timer errors.

Status: For the steppings affected, see the *Summary Table of Changes*.

11. **NB hang during stress and PCI Express link events**

Problem: An arbitration deadlock can occur when a PCI express device issues unaligned DMA writes to memory and on the same I/O unit, either a link goes down or a PCI Express-related internal MSI is generated.

Implication: Under these conditions, the NB may hang the system or silent data corruption may occur.

Workaround: A workaround for this issue which disables internally generated MSIs and adjusts arbiter timing is described in the BIOS Spec Update.

Status: For the steppings affected, see the *Summary Table of Changes*.

12. **NB may cause IERR# when using SYRE.CPURST with the Intel® Xeon™ Processor MP with 1 MB L2 Cache**

Problem: During CPU-only resets with Intel® Xeon™ processor MP with 1 MB L2 cache invoked via the NB's SYRE.CPURST bit, there is a window of 1/2 cycle where the NB will not be able to stop the CPU from driving a new FSB transaction before the NB is able to drive reset on the bus. If a transaction occurs during this window, the NB will queue a response, and drive it after the reset has been deasserted.

Implication: The processor will assert IERR# if such a spurious transaction is issued after reset.

Workaround: Programmers should not use the SYRE.CPURST bit with Intel Xeon processor MP with 1 MB L2 cache.

Status: For the steppings affected, see the *Summary Table of Changes*.

13. NB may trigger ECC errors during warm reset

- Problem:** The NB treats reset assertion as an asynchronous event. If the bus is active during a system reset, NB FSB I/O flip-flops may capture a cycle of data from the NB core as the core flip-flops are transitioning to the reset state. Because setup time is not met, some of the I/O flip-flops capture the pre-reset value and others capture the reset value.
- Implication:** The NB may drive spurious FSB data or exhibit a FSB strobe glitch during the reset sequence. The processor may detect an ECC or parity error in this case.
- Workaround:** Programmers should clear and ignore processor ECC, glitch, or parity errors logged immediately after a NB reset.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

14. NB PCI Express ports get stuck in reset state

- Problem:** The NB x8 PCI Express I/O units have a unique, narrow band of susceptibility when connected to an un-powered PCI Express end agents that have low impedances to ground. This low impedance causes the IOU to oscillate rapidly between Detect and Polling state.
- Implication:** If this oscillation continues for a long period (a few seconds) the unit may eventually become stuck in a reset state until the NB is power cycled or reset. Only x8 hot-plug slots with unpowered end agents are susceptible to this issue.
- Workaround:** A workaround for this issue which disables the NB PCIe transmitters when a slot is unpowered is described in the BIOS Spec Update.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

15. Receiver errors reported on NB when link is disabled

- Problem:** The NB is flagging the correctable PCI Express error “receiver error” when the link is disabled.
- Disabling any of the PCI Express ports by setting the LNKDIS bit, bit 4, in the EXP_LNKCTRL[7:1] register at offset 0x7C, device 1-7, function 0, will cause correctable receiver errors to be erroneously reported. The link errors are logged in bit 0 of register CORERRSTS[7:1] at offset 0x110, device 1-7, function 0. Bit 0 is the IO7ERR bit which denotes a correctable receiver error. This error cannot be cleared as long as the PCI Express port is held in the disabled state.
- Implication:** This issue affects all of the PCI Express ports regardless of link width configuration with exception of port B when it is configured as x8.
- Workaround:** Receiver error detection should be masked prior commanding a port to disable its link.

Workaround Implementation

Use the following procedure to disable a link:

1. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 4-7, function 0.
2. Set the disable link bit, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 1-7, function 0.

Use the following procedure to enable the link:

1. Clear the disable link, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 1-7, function 0.

2. If the port is A0 and this port meets conditions for erratum 7, “PCI Express Port A0 will not Train x1” then apply workaround for erratum 7, “PCI Express Port A0 will not Train x1”.
3. Wait 100 ms.
4. Clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 1-7, function 0.

Status: For the steppings affected, see the *Summary Table of Changes*.

16. Receiver errors reported on Port 4 when link is disabled in Port 6 in the 2 x8 configuration

Problem: The NB hub is reporting the correctable PCI Express error “receiver error” on port B when ports A0 and or A1 are disabled. This issue manifests itself only when port B is configured as a x8 link (ports B0 and B1 combined). This error can not be cleared as long as the link is held in the disabled state.

Disabling either PCI Express port A0 or A1 by setting the LNKDIS bit, bit 4, in the EXP_LNKCTRL[7:1] register at offset 0x7C, device 6-7, function 0, will cause correctable receiver errors to be erroneously reported on port B. The link errors are logged in bit 0 of register CORERRSTS[7:1] at offset 0x110, device 4, function 0. Bit 0 is the IO7ERR bit which denotes a correctable receiver error. This error cannot be cleared as long as the PCI Express port is held in the disabled state.

Implication: This issue affects only port B when it is configured as a x8 link (ports B0 and B1 combined). This issue manifests its self regardless of whether ports A0 and A1 are combined as a x8 link or are separate x4 links.

Workaround: Receiver error detection should be masked prior commanding a port A0 or A1 to disable its link, and should be left disabled as long as the link in port A0 or port A1 is disabled

Workaround Implementation

Use the following procedure to disable a link:

1. If port B is configured as a x8 link and either ports A0 and A1 are to be disabled.
2. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 6 or 7, function 0.
3. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 4, function 0.
4. Set the disable link bit, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 6 or 7, function 0.

Use the following procedure to enable the link:

1. If port B is configured as a x8 link and either ports A0 and A1 are to be enabled.
2. Clear the disable link, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 6 or 7, function 0.
3. If the port is A0 and this port meets conditions for erratum 7, “PCI Express Port A0 will not Train x1” then apply workaround for erratum 7, “PCI Express Port A0 will not Train x1”.
4. Wait 100 ms.
5. Clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 6 or 7, function 0.

6. If links on ports A0 and A1 are enabled; clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 4, function 0.

Status:

For the steppings affected, see the *Summary Table of Changes*.

Specification Changes

There are no Specification Changes in this revision of the Specification Update.

Specification Clarifications

1. Error detection for PCI Express devices not present should be disabled

To avoid spurious errors, BIOS should not enable error detection (via the UNCEDMASK register) for PCI Express devices which are not present or for slave devices (i.e. the upper x4 part of a x8 link). This detection is off by default.

2. PCI Express links enter compliance mode if connected to a powered down device

On a PCI Express link where the other device is powered down, the NB may detect the device, attempt to establish a link, fail, and then commence sending compliance patterns. This is a default behavior that occurs as per the PCI Express specification. When the device comes out of reset and tries to link train, the NB will stop transmitting the compliance pattern and try to link train. It is strongly recommended that BIOS set the EXP_CTRL.NO_COMPLIANCE (Devices 1-7, Function 0, Offset 4C, bit 0) to avoid one bad link from affecting other links during link training. This enables the link training state machine to exit Polling/Compliance and disabling any further entry into this state.

3. PCI Express: Configuration completion time-out enable needs non-configuration completion time-out to be enabled in order to work

The PCI Express configuration completion time-out timer requires both of the TIMEOUT_ENABLE and TIMEOUT_ENABLE_CFG bits to be set. Simply setting the TIMEOUT_ENABLE_CFG bit will not enable the completion time-out timer. These bits are found in the EXP_CTRL[7:1] register.

This affects all of the PCI Express ports. There are only three choices of completion time-outs:

1. No completion time-outs.
2. Non-configuration completion time-outs.
3. Both configuration and non-configuration time-outs.

There is no way to enable just the configuration time-out counter.

Use of configuration completion time-outs is not recommended (the PCIE specification requires that all non-posted transactions be timed-out excepting configuration cycles which should not time-out), but if you require them, use the following workaround. When time-out for configuration transactions is needed, both of the TIMEOUT_ENABLE and TIMEOUT_ENABLE_CFG bits in the EXP_CTRL[7:1] register need to be set.

Process Implementation

To enable both configuration and non-configuration transaction time-outs:

1. Set the TIMEOUT_ENABLE_CFG bit, bit 23 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.
2. Set the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

To enable just non-configuration transactions time-out:

1. Set the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

To disable both configuration and non- configuration transactions time-outs:

1. Clear the TIMEOUT_ENABLE_CFG bit, bit 23 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.
2. Clear the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

Documentation Changes

There are no Documentation Changes in this revision of the Specification Update.

§